

Fragment Approach to Constrained Density Functional Theory Calculations using Daubechies Wavelets

Laura E. Ratcliff,^{1,2,*} Luigi Genovese,² Stephan Mohr,² and Thierry Deutsch²

¹*Argonne Leadership Computing Facility, Argonne National Laboratory, Illinois 60439, USA*

²*Univ. Grenoble Alpes, CEA, INAC-SP2M, L_Sim, F-38000, Grenoble, France*

(Dated: March 29, 2015)

In a recent paper we presented a linear scaling Kohn-Sham density functional theory (DFT) code based on Daubechies wavelets, where a minimal set of localized support functions is optimized *in situ* and therefore adapted to the chemical properties of the molecular system. Thanks to the systematically controllable accuracy of the underlying basis set, this approach is able to provide an optimal contracted basis for a given system: accuracies for ground state energies and atomic forces are of the same quality as an uncontracted, cubic scaling approach. This basis set offers, by construction, a natural subset where the density matrix of the system can be projected. In this paper we demonstrate the flexibility of this minimal basis formalism in providing a basis set that can be reused *as-is*, i.e. without reoptimization, for charge-constrained DFT calculations within a *fragment* approach. Support functions, represented in the underlying wavelet grid, of the template fragments are roto-translated with high numerical precision to the required positions and used as projectors for the charge weight function. We demonstrate the interest of this approach to express highly precise and efficient calculations for preparing diabatic states and for the computational setup of systems in complex environments.

I. INTRODUCTION

Density functional theory (DFT)^{1,2} is arguably the most popular approach to electronic structure calculations for a wide range of systems. However, it suffers from various well-known limitations, like for example the self-interaction problem^{3,4} which can result in electron delocalization errors, and the fact that it is in principle a ground state theory only. For these reasons, the DFT formalism has been extended in the form of constrained DFT (CDFT)⁵ to include an additional constraint on the density, so that the lowest energy state *satisfying a given condition* can instead be found. When a reasonable guess for such a condition is at hand, it can therefore be used both to find a particular excited state of the system and to localize the electronic density in such a way as to prevent spurious delocalization, and thus provides a way of overcoming the above problems. Of course, time-dependent DFT (TDDFT)⁶ can be used to find multiple excited states, however when one is interested in a particular excited state CDFT can be advantageous, especially given that the additional costs associated with adding a constraint are relatively low. Furthermore, TDDFT also suffers from self interaction problems, and can give inaccurate results for certain types of excited states, including charge transfer excitations.

Constrained DFT has been implemented in a number of codes, using both localized basis sets⁷ and plane waves^{8,9}, and has been successfully applied in a variety of contexts, including charge constrained molecular dynamics⁸, the calculation of the correct energy alignment of metal/molecule interfaces¹⁰ and the calculation of electronic coupling matrix elements^{11,12}. For a general overview of CDFT see Ref. 13.

As with all DFT calculations, the choice of basis set has a large impact on both the accuracy and computa-

tional cost of CDFT. One way of accessing large systems is to reformulate the standard cubic scaling approach to DFT in terms of localized orbitals, or ‘support functions’, which we will discuss further in the following section. We wish to perform CDFT calculations on large systems using such an approach, whilst maintaining the high accuracy associated with systematic basis sets. As such, we require a basis set which is at the same time localized and systematic. For this reason, we have chosen to use a Daubechies wavelet basis set¹⁴, as it is a systematic basis set exhibiting the desired properties of compact support in both real and Fourier space and can be chosen to be orthogonal. Wavelet basis sets have an inherent flexibility, in that they allow for multiresolution grids, which is particularly useful for inhomogeneous systems. Combined with the ability to explicitly treat charged systems in open boundary conditions, wavelets provide an ideal basis set for accurate CDFT calculations of large systems.

In this paper, we show that the combination of a support function approach with a wavelet basis set allows for the definition of a flexible fragment based approach to CDFT, which can further reduce the computational cost, particularly for very large systems. In this approach, a set of support functions are optimized for an isolated (small) molecule, or ‘fragment’, and reused as a fixed basis in a larger system containing many of these molecules e.g. a solvent. It is then straightforward to associate the constrained charge with a given fragment using a Löwdin like definition of the CDFT weight function. However, in the larger system each molecule may well have a different orientation and so the support functions, which are described in terms of the fixed wavelet grid, cannot simply be duplicated for each molecule.

Therefore, we have developed a scheme to reformat the support functions for arbitrary roto-translations using interpolating scaling functions. This interpolation, thanks




FIG. 1. Least asymmetric Daubechies wavelet family of order $2m = 16$; both the scaling function $\phi(x)$ and wavelet $\psi(x)$ differ from zero only within the interval $[1 - m, m]$.

to the properties of the underlying basis set, results in only a negligible loss of accuracy and so the support functions can be directly reused, reducing the computational cost by an order of magnitude compared to optimizing the support functions from scratch for the full system.

In the following sections we will first summarize our approach to large scale DFT calculations using localized support functions represented in a wavelet basis set¹⁵, as implemented in the BigDFT electronic structure code¹⁶. We will then outline our implementation of CDFT, following which we will explain our fragment approach, including a description of the reformatting scheme, validating our method with calculations on prototypical systems. Finally, we will present an application of CDFT for the fullerene C_{60} in two different environments, through which we will demonstrate the flexibility and potential of a fragment based approach.

II. METHODOLOGY

A. Linear scaling DFT with wavelets

We and others have recently presented a newly developed method for DFT calculations on large systems, which combines the use of a minimal localized basis of ‘support functions’ with the use of an underlying wavelet basis set¹⁵. This method has been implemented in BigDFT, which uses the orthogonal least asymmetric Daubechies family of order 16, which are depicted in Fig. 1. The Kohn-Sham (KS) orbitals are expressed in terms of the support functions via a set of coefficients c_i^α :

$$|\Psi_i\rangle = \sum_{\alpha} c_i^\alpha |\phi_\alpha\rangle, \quad (1)$$

where the support functions are represented directly in the wavelet basis set localized on a 3 dimensional grid, so that they can be thought of as adaptively contracted wavelets. Rather than working directly with the KS orbitals, we instead work in terms of the density matrix, $\rho(\mathbf{r}, \mathbf{r}')$, which is itself defined in terms of the support functions and the density kernel, $K^{\alpha\beta}$:

$$\rho(\mathbf{r}, \mathbf{r}') = \sum_{\alpha, \beta} \phi_\alpha(\mathbf{r}) K^{\alpha\beta} \phi_\beta(\mathbf{r}'). \quad (2)$$

The density matrix has been shown to decay exponentially with distance for systems with a gap thanks to the so-called nearsightedness principle^{17–22}, and thus a formulation in terms of the density matrix allows us to take advantage of this to achieve linear scaling with the number of atoms in the system, thereby avoiding the cubic scaling of standard approaches to DFT. From this the charge density is calculated directly from the support functions and density kernel. Similarly, the band structure energy and the charge of the system can be calculated from the density kernel via:

$$E_{\text{BS}} = \text{Tr} [\mathbf{KH}], \quad N = \text{Tr} [\mathbf{KS}] \quad (3)$$

where \mathbf{H} indicates the Hamiltonian matrix in the basis of the support functions, and \mathbf{S} is the support function overlap matrix.

The support function formalism allows one to *map* the degrees of freedom of KS orbitals into a *localized* description, that can be directly put in relation with atomic positions. In practice, the support functions are truncated within spherical localization regions with a user-defined radius, and some additional truncation must be applied to the density kernel which is then exploited via sparse matrix algebra to achieve a fully linear scaling algorithm. Therefore, to some extent, a given support function ϕ_α can be associated to the atom a where its localization region is centered. In order to achieve accurate results, both the support functions and density kernel are optimized during the calculation, that is the energy is minimized with respect to both quantities. Providing the localization regions are sufficiently large, this results in a minimal localized basis set with an accuracy equivalent to the underlying basis set.

The general scheme is common to other basis optimization for density-matrix minimization based linear scaling DFT codes, e.g. ONETEP²³ and Conquest²⁴, with the addition of a few novel features. These include the application of a confining potential to the KS Hamiltonian, which ensures the support functions remain localized. In order to apply the confining potential consistently, we also enforce an approximate orthogonality constraint on the support functions, in contrast with other approaches which use fully non-orthogonal support functions^{23,24}. Furthermore, the properties of the description in the wavelet basis are such that the algorithm guarantees that the Pulay contribution to the atomic forces can safely be neglected, so the forces can be calculated accurately and cheaply.

The method can be divided into two key components: the optimization of the support functions (either with or without a confining potential), and the optimization of the density kernel. This latter point can be achieved via a choice of schemes, as detailed previously¹⁵. These include a direct minimization approach, where the coefficients c_i^α are first updated using DIIS or steepest descents to minimize the band structure energy and then used to construct the density kernel, and the Fermi Operator Expansion (FOE) method, where the density kernel is ex-

pressed as a function of the Hamiltonian matrix which can be evaluated numerically using a Chebyshev polynomial expansion. Close attention has also been paid to the parallelization of the code, such that massively parallel machines can be exploited to perform large scale calculations (see also Ref.²⁵). For charged calculations, we have found the use of the direct minimization method to be the most suitable, due to a reduction in the occurrence of charge sloshing during convergence, and the flexibility afforded by working with the wavefunction coefficients rather than directly with the density kernel as in the FOE method.

B. Atomic charge analysis

The mapping between electronic and localized degrees of freedom which is provided by the support function formalism allows one to perform an accurate atomic charge analysis, meaning that each atom is assigned a partial net charge, such that the electrostatic properties of the system are conserved. Obviously this conservation is only possible within a certain limit, as one is mapping a continuous quantity (the electronic charge) to a discrete quantity (the atomic point charges). If, however, the error introduced by this mapping onto point charges is small enough, the system under investigation can be reasonably approximated by a simple setup of charged point particles, which paves the way for future applications such as coupling different levels of accuracy within the same calculation.

Given the overlap matrix \mathbf{S} and the density kernel \mathbf{K} , the partial charge located on atom a can be defined by the so-called Löwdin charge:

$$q_a = \sum_{\alpha'}^{(a)} \left(\mathbf{S}^{1/2} \mathbf{K} \mathbf{S}^{1/2} \right)_{\alpha' \alpha'}, \quad (4)$$

where the sum runs over all support functions α' which are located on atom a . Obviously $\sum_a q_a = \text{tr}(\mathbf{K}\mathbf{S}) = N$, i.e. the total charge (the monopole) is conserved. In order to check whether higher multipoles can also be conserved, we compared the dipole moment calculated using this approach with that calculated using the continuous charge density. In addition a comparison with the dipole moment calculated with the cubic scaling version of BigDFT was done as a reference. The values for a strongly polarized molecule (H_2O) and a non-polarized one (C_{60}) are given in Tab. I.

As a second test of the reliability of our method we directly compared the atomic point charges with those calculated by performing a Bader charge analysis of the charge density calculated using the cubic scaling approach. As can be seen from Tab. II, the differences between the exact results are smaller with our approach. In particular, for the C_{60} fullerene, reasons of symmetry impose that the charge should be equally distributed and no atom should carry a net charge. As can be seen, the

Löwdin procedure comes closer to this result than the Bader analysis.

C. Constrained DFT

The general idea of constrained DFT is to force a charge to remain localized in a given region of the simulation space. This is achieved via the addition of a Lagrange multiplier term to the Kohn-Sham energy functional which enforces a given constraint on the resulting electronic density, so that rather than being the ground-state density of the system, the density instead corresponds to a particular excited state. This Lagrange multiplier can also be thought of as an additional applied potential, otherwise referred to as the constraining potential. The constraint can also take a number of other forms, but for the purposes of this work we are interested only in constraining the charge. The new functional therefore becomes:

$$W[\rho, V_c] = E_{\text{KS}}[\rho] + V_c \left(\int w_c(\mathbf{r}) \rho(\mathbf{r}) d\mathbf{r} - N_c \right), \quad (5)$$

where E_{KS} is the Kohn-Sham energy functional, V_c is the aforementioned Lagrange multiplier, N_c is the required charge within the specified region and $w_c(\mathbf{r})$ is a weight function which defines this region. The weight function and N_c are defined in advance, however the value of V_c which correctly enforces the constraint must be found during the calculation. Wu and Van Voorhis²⁶ demonstrated that the lowest energy state for which the constraint is correctly applied is in fact a maximum with respect to V_c and so it becomes possible to efficiently determine the correct V_c . It is also straightforward to add multiple constraints to the system, and indeed one is frequently interested in constraining the charge difference between two regions. This feature is important for the simulation of charge-transfer excitations within the CDFT formalism.

Rewriting the new functional (Eq. 5) in density matrix form as⁷:

$$W[\rho, V_c] = E_{\text{KS}}[\rho] + V_c (\text{Tr}[\mathbf{K}\mathbf{w}_c] - N_c), \quad (6)$$

the charge constraint is easily added to the existing algorithm in BigDFT. This construction requires the weight matrix, $w_{\alpha\beta}$, which is defined via the weight function as

$$w_{\alpha\beta} = \int \phi_\alpha(\mathbf{r}) w_c(\mathbf{r}) \phi_\beta(\mathbf{r}) d\mathbf{r}. \quad (7)$$

It then remains to define the weight function, for which a number of different schemes exist. The support function approach used here lends itself to a Löwdin like definition, which is analogous to that used above to determine atomic charges. Using this approach we directly construct the weight matrix via

$$w_{\alpha\beta} = \left(\mathbf{S}^{\frac{1}{2}} \mathbf{P} \mathbf{S}^{\frac{1}{2}} \right)_{\alpha\beta}, \quad (8)$$

	cubic	linear	
	exact dipole	exact dipole	point charge approx.
H ₂ O	(0.463, -0.506, -0.186)	(0.466, -0.510, -0.187)	(0.606, -0.668, -0.247)
norm	0.711	0.716	0.935
$d_{ex} \cdot d$	1	0.9999996	0.9999897
C ₆₀	(-0.0004, -0.0004, -0.0004)	(-0.025, -0.025, -0.025)	(-0.055, -0.055, -0.055)

TABLE I. Dipole moments calculated using the exact charge density for the cubic and linear scaling approaches, respectively, and using the partial atomic point charges. All values are given in atomic units.

	cubic – Bader	linear – Löwdin
H ₂ O	(-1.27, 0.61, 0.67)	(-0.83, 0.42, 0.41)
C ₆₀	0.061 ± 0.045	0.003 ± 0.002

TABLE II. Atomic point charges, calculated by a Bader analysis of the charge density from the cubic scaling approach, and the Löwdin procedure using the density kernel and overlap matrix from the linear scaling approach. For H₂O we indicate the values on all three atoms, for C₆₀ we give the mean of the absolute values together with the standard deviation. All values are given in atomic units.

where \mathbf{S} is the overlap matrix between support functions and \mathbf{P} is a projection matrix, defined as 1 for all support functions belonging to the region where a constraint is being applied, and 0 elsewhere. Alternatively, if one is constraining a charge difference between two regions, it should be set to 1 on one of the regions, -1 on the other, and 0 elsewhere.

The final step, once the weight function has been defined, is to derive a scheme for finding the correct value of V_c for a given charge constraint value, N_c . There are two possible approaches to the optimization. In the first approach, one can find the optimum value of V_c at each step of the self-consistent density optimization, i.e. the ground state density is updated in an outer loop, with V_c updated in an inner loop. Alternatively, the second approach consists of fully minimizing the functional $W[\rho, V_c]$ of Eq. 6 with respect to the density for a fixed value of V_c , updating V_c and finding the new minimum density, then repeating to convergence, i.e. the maximization with respect to V_c is performed in an outer loop with the ground state density found self-consistently in an inner loop. We chose the latter, as it was observed to be more stable. We use Newton’s method to update V_c , with the second derivative calculated using a finite difference approach.

D. Fragment approach

The combination of the novel features described above and the use of a wavelet basis set make this approach ideal for the application of CDFT to large systems. In particular the ability to reuse the support functions can result in significant savings for e.g. geometry optimizations and calculations on charged systems, as previously




FIG. 2. The fragment approach as illustrated for a cluster of water molecules: the support functions are initially optimized for an isolated water molecule and then duplicated for a collection of water molecules, avoiding the need for optimization in the larger system.

demonstrated¹⁵. Furthermore, this idea of support function reuse can be extended to a fragment based approach, which is similar to the so-called fragment orbital method which has been used to calculate electronic coupling matrix elements^{12,27,28}.

The central idea is to take a group of atoms, or more specifically an isolated molecule, and fully optimize the support functions. These support functions are then used as a fixed basis for a system containing several molecules, as illustrated in Fig. 2 for a simple example. We refer to the initial molecule for which the support functions were optimized as the ‘template’ molecule.

As the support functions are kept fixed in the fragment approach, $w_{\alpha\beta}$ need only be calculated once at the start of the calculation, after which it remains fixed. Furthermore, due to the quasi-orthogonality of the support functions, when the fragment approximation is justified

$\mathbf{S}^{\frac{1}{2}}$ can in general be calculated using a Taylor approximation, and so the calculation of the weight matrix adds very little overhead to the calculation.

In this work we focus on systems where the respective fragments are well defined, and thus the support functions generated from the isolated fragments can be used for the full system with a minimal impact on the accuracy. In cases where electrons are being added to a fragment, it is important to ensure that the lowest unoccupied molecular orbital (LUMO) and, if necessary, the next few states in energy are sufficiently well represented by the support function basis. As discussed elsewhere¹⁵, this can be achieved using the direct minimization formalism to optimize a few additional states during the isolated calculation without adding a charge to the system. For systems where the fragments are less well defined the implementation could in principle be extended to further optimize the support functions for the combined system, either in a neutral state or while the charge constraint is being enforced. In such cases, the Löwdin approach is also expected to be less accurate, and so it would be desirable to use an alternative form for the weight function.

Finally, it should be mentioned that there are some subtleties related to the initial guess for the charge density. This depends on the initial density kernel, which is constructed from the fragment KS orbitals. For neutral calculations it is straightforward to use the fragment orbitals and occupancies directly from the isolated calculations, however for charged calculations some additional input is required. One approach would be to occupy the fragment orbitals in order of their energies, however this can lead to charge distributions which are significantly different from the required constraint. This can result in slow convergence, or even worse, problems with local minima. A better approach should therefore take into account the effect of the constraining potential on the fragment orbital energies. This can be done by assigning occupation numbers so that any excess/deficit in charge is localized on the same fragment as the constraint, so that the initial density already satisfies the charge constraint. Alternatively, the risk of encountering a local minimum can be reduced by adding a degree of noise to the fragment orbitals, or by completely randomizing the initial guess, subject to the correct overall charge. However, such an approach is in general much slower to converge, and thus the latter strategy is generally not recommended.

E. Reformatting scheme for roto-translations

As the support functions are defined in terms of an underlying grid of wavelets, in order to implement a fragment approach it is necessary to have a scheme for reformatting the support functions due to a change in atomic positions. We are frequently interested in situations where a molecule has been rotated and translated, for example when calculating electronic coupling matrix

elements in a dimer for varying angles between the two monomers. Therefore, we have developed and implemented a scheme for reformatting the support functions, given the axis and angle of rotation between some initial and final positions for a given fragment mass center.

A fragment of the system is defined by the user via a list of atomic positions, which should of course be in bijection with the atom list defining the template fragment. Therefore the first problem is to identify the combination of translation and rotation which sends the template fragment to the position of the system's fragment. As a first step, two reference systems are chosen such that the fragment center of mass is in the same position. This operation is equivalent to finding the translation between the template and the system. We then have two lists of atomic positions, $\{\mathbf{R}_a^{T,S}\}$, where T and S label the template and system fragment respectively, and the subscript a , indicating the atom, ranges from 1 to N , the number of atoms in the fragment.

If the system fragment is a *rigid* displacement of the template (i.e. its internal coordinates are unchanged), the rotation matrix \mathcal{R} we seek is such that $\mathbf{R}_a^S = \sum_{b=1}^N \mathcal{R}_{ab} \mathbf{R}_b^T$. In general, we should assume there is a slight modification of the internal coordinates, as the geometry of the fragment might be affected by the interaction with the environment. In this case, the matrix \mathcal{R} is such as to *minimize* the cost function

$$J(\mathcal{R}) = \frac{1}{2} \sum_{a=1}^N \|\mathbf{R}_a^S - \sum_{b=1}^N \mathcal{R}_{ab} \mathbf{R}_b^T\|^2. \quad (9)$$

The determination of the matrix \mathcal{R} in such a manner constitutes a version of the well-known Kabsch algorithm²⁹ (also known as Wahba's problem)³⁰, which can be solved by the Singular-Value Decomposition of the 3-by-3 matrix³¹ $\mathcal{B}_{ij} = \sum_{a=1}^N (\mathbf{R}_a^S)_i (\mathbf{R}_a^T)_j$. After having found two matrices \mathcal{U} and \mathcal{V} and a diagonal matrix \mathcal{S} such that $\mathcal{B} = \mathcal{U} \mathcal{S} \mathcal{V}^t$, the optimal rotation is

$$\mathcal{R} = \mathcal{U} \mathcal{D} \mathcal{V}^t \quad \mathcal{D} = \text{diag}(1, 1, \det(\mathcal{U}) \det(\mathcal{V})) \quad (10)$$

The value of $J(\mathcal{R})$ defined in (9) might then be used to quantify the validity of the rigid transformation approach. In the case where its value is below a given threshold (fixed to 10^{-3} in our case), we may proceed with the reformatting of the template basis functions, which will be denoted by $|\phi_\alpha^T\rangle$ in what follows.

As described in Refs.^{15,16,32}, from the expression of $|\phi_\alpha^T\rangle$ in a Daubechies wavelets basis set, the so-called "magic-filter" transformation can be used to define a real space representation of the basis functions, given in terms of one-dimensional interpolating scaling functions (ISF)

$$\phi_\alpha^T(x, y, z) = \sum_{i,j,k} c_{ijk} \varphi_i(x) \varphi_j(y) \varphi_k(z), \quad (11)$$

where $\varphi_i(x) \equiv \varphi(x/h - i)$ is one element of the ISF basis set, which is constituted of uniform translations of

the mother function $\varphi(t)$ over the points of a uniform grid of spacing h , covering the entire simulation domain. These points are labeled by indices ijk . The points $\mathbf{r}_{ijk} = (hi, hj, hk)$ therefore lie within the box containing the support of $\phi_\alpha^T(\mathbf{r})$.

This real-space expression is optimal in the sense that it preserves the same moments of the original representation given in Daubechies wavelets. The interpolating property of the ISF basis set is such that $c_{ijk} = \phi_\alpha^T(\mathbf{r}_{ijk})$. Let us suppose we have a one-dimensional function expressed in ISF, namely $f(x) = \sum_i f_i \varphi_i(x)$. We know that $f_i = f(hi)$. If we want to translate the function f by a displacement Δ , and express this function in the ISF basis, we have $f(x + \Delta) = \sum_i f'_i \varphi_i(x)$, with

$$f'_i = f(hi + \Delta) = \sum_j f_{i-j} t_j^\Delta, \quad (12)$$

where the filter $t_j^\Delta = \varphi(j + \Delta/h)$ implements the (uniform) translation. This filter has a limited extension (the same as the function $\varphi(x)$) and of course $t_j^{hk} = \delta_{j,-k}$.

Imagine now we have a *different* ISF basis set $\{\varphi_I(\tilde{x})\}$ defined on a uniform grid spacing of separation \tilde{h} and a reference frame $\tilde{x}_I = \tilde{h}I$, which is related to x by a more complicated transformation $\tilde{x}(x)$ of the coordinate space. If this transformation can be inverted, by $x(\tilde{x})$, then a new function $\tilde{f}(\tilde{x}) \equiv f(x(\tilde{x}))$ can be defined in this frame. For each grid point I , it is then always possible to find \bar{i} in the old frame such as to minimize the absolute value of $\Delta_I \equiv x(\tilde{x}_I) - h\bar{i}$.

Using the above relations we might approximate $\tilde{f}(\tilde{x}) \simeq \sum_I \tilde{f}_I \varphi_I(\tilde{x})$, where

$$\tilde{f}_I = \tilde{f}(\tilde{x}_I) = f(h\bar{i} + \Delta_I) = \sum_j f_{\bar{i}-j} t_j^{\Delta_I}. \quad (13)$$

If the transformation \tilde{x} is a continuous function of x which varies slowly enough, this is in general a rather good approximation (see Fig. 3).

This framework can be easily generalized to a roto-translation in three dimensions. Indeed, we would like to estimate the function

$$\phi_\alpha^S(\tilde{\mathbf{r}}) \equiv \phi_\alpha^T(\mathbf{r}(\tilde{\mathbf{r}})) \simeq \sum_{I,J,K} \tilde{c}_{IJK} \varphi_I(\tilde{x}) \varphi_J(\tilde{y}) \varphi_K(\tilde{z}), \quad (14)$$

where the coordinates $\tilde{\mathbf{r}} = (\tilde{x}, \tilde{y}, \tilde{z})$ are defined as

$$\tilde{\mathbf{r}} = \mathcal{R} \cdot \mathbf{r} \quad (15)$$

where \mathcal{R} is calculated by Eq. (10). In addition a rigid shift vector $\mathbf{s} = (s_x, s_y, s_z)$ is defined as the difference between the coordinates of the center of mass of the two fragments. If the rotation is the identity matrix, the template reference frame is then $\tilde{\mathbf{r}} = \mathbf{r} + \mathbf{s}$. As in the one dimensional case presented above, the interpolation depends on the *inverse* mapping $\mathbf{r}(\tilde{\mathbf{r}})$. We detail in the following a procedure to identify such a function.

The coefficients \tilde{c}_{IJK} of $\phi_\alpha^S(\tilde{\mathbf{r}})$ can be found in three steps. We first start by considering the transformation

law for \tilde{x} . This transformation can be thought of as a function of the template coordinates \mathbf{r} :

$$\tilde{x}(x, y, z) = \mathcal{R}_{11}x + \mathcal{R}_{12}y + \mathcal{R}_{13}z. \quad (16)$$

In the same spirit as Eq. (13), we may invert Eq.(16) with respect to one template coordinate $t = x, y, z$ into \tilde{x} in the system's reference frame. The choice of the variable t depends on the entries of the rotation matrix, and it is in general given by the coordinate which is multiplied by the coefficient of the highest absolute value in Eq. (16). This choice guarantees that the t variable is the one for which $\tilde{x} - t$ is slowly varying. Let us imagine $t = x$ for this example. We can define the function

$$\begin{aligned} \phi_\alpha^{(1)}(\tilde{x}, y, z) &= \phi_\alpha^T(x(\tilde{x}, y, z) - s_x, y, z) \\ &= \sum_{I,j,k} \tilde{c}_{I,j,k} \varphi_I(\tilde{x}) \varphi_j(y) \varphi_k(z), \end{aligned} \quad (17)$$

by proceeding for all j, k , as described in Eq.(13), to define the coefficients $\tilde{c}_{I,j,k}$. The second step is related to the expression of \tilde{y} . Depending on the choice of the variable t in the first step, we have to consider one of these three relations:

$$\mathcal{R}_{11}\tilde{y} = \mathcal{R}_{21}\tilde{x} + \mathcal{R}_{33}y - \mathcal{R}_{32}z, \quad (18)$$

$$\mathcal{R}_{12}\tilde{y} = \mathcal{R}_{22}\tilde{x} - \mathcal{R}_{33}x + \mathcal{R}_{31}z, \quad (19)$$

$$\mathcal{R}_{13}\tilde{y} = \mathcal{R}_{23}\tilde{x} + \mathcal{R}_{32}x - \mathcal{R}_{31}y, \quad (20)$$

which hold when in the first step $t = x, y, z$ respectively. These relations can be derived using the orthogonality of the rotation matrix \mathcal{R} . This function can now be inverted with respect to one of the old variables. Again, this choice will depend on the values of the coefficients multiplying each variable.

In our example, we have to consider the relation (18) as we have chosen $t = x$ in the first step. We choose to invert the relation with respect to z , having $z = z(\tilde{x}, \tilde{y}, y)$. In this case we will have, as a second step

$$\begin{aligned} \phi_\alpha^{(2)}(\tilde{x}, \tilde{y}, y) &= \phi_\alpha^{(1)}(\tilde{x}, y, z(\tilde{x}, \tilde{y}, y) - s_z) \\ &= \sum_{I,j,J} \tilde{c}_{I,j,J} \varphi_I(\tilde{x}) \varphi_j(\tilde{y}) \varphi_J(y). \end{aligned} \quad (21)$$

In the third step, the remaining variable, (which is y for the illustrated example), can be directly obtained from the inverse relation

$$\mathbf{r} = \mathcal{R}^{-1} \cdot \tilde{\mathbf{r}} = \mathcal{R}^t \cdot \tilde{\mathbf{r}}, \quad (22)$$

which is easier to express as \mathcal{R} is an orthogonal matrix. In our case, the final result is therefore

$$\phi_\alpha^S(\tilde{x}, \tilde{y}, \tilde{z}) = \phi_\alpha^{(2)}(\tilde{x}, y(\tilde{x}, \tilde{y}, \tilde{z}) - s_y, \tilde{z}). \quad (23)$$

We recall that the definition of $\phi_\alpha^{(1,2)}$ depends on the order of the operation. Here we have chosen to interpolate first with respect to x , then z and y . The best choice of order depends only on the entries of the matrix \mathcal{R} .

1. Accuracy

In order to assess the accuracy of the reformatting scheme, we have applied it to a water molecule undergoing a series of rotations. Support functions were generated for a template water molecule, using a dense grid with a spacing of 0.132 Å, and were then reused for water molecules in a variety of different orientations using a less dense grid with a spacing of 0.185 Å. As a point of comparison, calculations were also performed for each orientation fully optimizing the support functions with a grid spacing of 0.185 Å. This allows us to quantify both the error introduced by the support function reformatting and the errors due to representing the wavefunctions on a fixed grid, i.e. the so-called ‘eggbox effect’. The eggbox effect of the standard cubic scaling approach is also presented. The computational setup has been chosen such that the difference in ground state energies between the cubic and support function approaches is of the order of 1 meV/atom.

The results are shown in Fig. 3, where we can see that the eggbox effect is of the order of 0.1 meV/atom. As both the cubic and linear scaling approaches use the same underlying grid, the variation is similar in each case. The error due to the interpolation also remains small – less than a few meV/atom. Importantly, the overall error for the reformatted calculations remains of the same order of magnitude as that due to the selected localization radii of the support functions.

III. RESULTS

Below we present results for three different systems, where the first two can be validated against CDFT implementations in other codes. In each case we use the local density approximation (LDA) exchange-correlation functional³³ and HGH pseudopotentials³⁴ within isolated boundary conditions. For carbon, nitrogen and oxygen we use four support functions per atom, for hydrogen we use one, and for zinc we use nine.

A. N₂

Wu and Van Voorhis have previously studied N₂³⁵ and so this system provides a useful test case. We used a grid spacing of 0.185 Å with support function radii of 7.4 Å, i.e. completely filling the simulation cell; here we aim to validate only the general correctness of the implementation of CDFT rather than the full fragment approach. Fixing the bond length at 1.12 Å, we have varied the charge separation between the two atoms, results for which are shown in Fig. 4. For simplicity, the calculations were performed only in the spin-averaged sense. Our results are closer to those obtained by Wu and Van Voorhis using a Becke weight population than the Löwdin scheme, however given that we have used the LDA whereas they

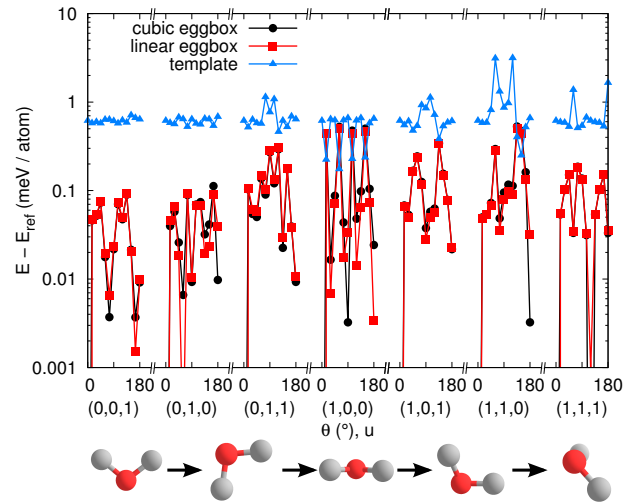


FIG. 3. Plot showing the energy variation for a water molecule rotated through different angles (θ) and axes of rotation (u). Results are shown for the standard cubic scaling approach (‘cubic eggbox’), fully optimized support functions (‘linear eggbox’) and a fixed support function basis generated for a template molecule (‘template’). The cubic reference is the energy at the initial orientation calculated using the cubic approach, for the linear and template approaches it is the same quantity calculated in the fully optimized support function basis. There is a roughly constant error of 1 meV/atom in the support function basis compared to the cubic scaling approach. Selected orientations are shown along the bottom.




FIG. 4. Change in energy with respect to an unconstrained calculation and applied potential value for differing charge separations in N₂.

used B3LYP we do not expect exact agreement. Furthermore, we should recall that the fragment approach presented here is aimed at systems where the donor and acceptor are well separated, whereas the use of support functions optimized for an isolated nitrogen atom is necessarily an approximation in this case. Nonetheless, we have successfully reproduced the correct trends for both the energy and the Lagrange multiplier.




FIG. 5. The ZnBC-BC model complex.

B. ZnBC

As a more elaborate test case we take the zincbacteriochlorin-bacteriochlorin (ZnBC-BC) complex, which has also been studied previously in some detail, both with CDFT^{26,35} and other approaches, e.g. Refs. 36 and 37. This system is ideally suited to our approach as the donor and acceptor are clearly separated. Furthermore, TDDFT has been shown to give incorrect energies for the $\text{ZnBC}^+-\text{BC}^-$ and $\text{ZnBC}^--\text{BC}^+$ charge transfer (CT) excited states³⁷ and so the advantages of CDFT are clear. It has previously been demonstrated that the differences between a (1,4)-phenylene-linked ZnBC-BC complex and a model complex where the link is eliminated are small³⁷; for simplicity we therefore choose to use the latter, where the distance between the two previously linked carbon atoms is 5.84 Å, as depicted in Fig. 5. Taking the coordinates from Ref. 37, we relaxed the isolated ZnBC and BC molecules separately, then built the model complex without further relaxation. We used a grid spacing of 0.185 Å and localization radii of 5.82 Å. To assess the accuracy of the fragment support functions we compare the neutral energies for the model complex with those obtained using cubic scaling BigDFT. The results are shown in Tab. III, where we can see that the error for both the model complex and the isolated molecules is less than 1 meV/atom.

The energies for the two CT excited states relative to the unconstrained DFT ground state are 3.71 eV for $\text{ZnBC}^+-\text{BC}^-$ and 3.98 eV for $\text{ZnBC}^--\text{BC}^+$, which is consistent with previous results^{26,35}. We can also gain some insight into the nature of these CT states by plotting the difference in the electronic density between the neutral and constrained calculations, as in Fig. 6. Not only is the charge transfer characteristic clear, the plot for $\text{ZnBC}^+-\text{BC}^-$ also shows remarkably good agreement with previous calculations that used the significantly more expensive Bethe-Salpeter approach³⁶, which confirms that CDFT can be used to obtain physically relevant CT excitons, and provide a reliable estimation of the corresponding excitation energies.

We have also plotted the relationship between the constraining potential, V_c , the total energy relative to the unconstrained calculation, ΔE , and the charge difference between the two molecules, N_c . This is shown in Fig. 7,

	cubic	frag.	diff.
	(eV)	(eV)	(meV)
ZnBC	-4472.626	-4472.604	22.6
BC	-4471.998	-4471.980	17.7
ZnBC-BC	-8944.629	-8944.575	54.1
$\text{ZnBC}^--\text{BC}^+$	-	-8940.592	-
$\text{ZnBC}^+-\text{BC}^-$	-	-8940.860	-

TABLE III. Energies for isolated ZnBC and BC, the neutral model ZnBC-BC complex and the two lowest energy CT states, as calculated using standard BigDFT ('cubic') and the fragment approach ('frag.'). Where applicable the difference between the two approaches is also indicated ('diff.').




FIG. 6. Density differences between the neutral and charged calculations for the two charge transfer states. Red (blue) indicates an increase (decrease) in the electronic charge density with respect to the neutral.

where $N_c = 1$ corresponds to $\text{ZnBC}^+-\text{BC}^-$ and $N_c = -1$ corresponds to $\text{ZnBC}^--\text{BC}^+$; our results agree well with previous calculations^{26,35}, despite the use of a different exchange-correlation functional. This test highlights the robustness of the method – in order for the correct value of V_c to be found within a minimal number of iterations of the constraint loop, there should be a smooth relationship between a given V_c and the resulting N_c . If for certain values of V_c the convergence is insufficient, such that the final charge deviates from the correct value, this will negatively impact the search for the correct V_c . We observed that in general such a smooth curve is straightforward to obtain, given a reasonable initial guess for the density kernel and therefore charge density. As discussed in Section II C, this can be achieved by defining the ini-




FIG. 7. Applied potential value and change in energy compared to an unconstrained calculation for differing charge separations in the ZnBC-BC model complex.

tial occupancies in a manner which is consistent with the desired charge difference.

C. C_{60}

In order to accurately calculate material properties it is important to account for environmental effects, e.g. by including a solvent or neighboring molecules in a molecular material. However, this can considerably increase the cost of a simulation, as in the case of large systems in solution where the solvent must fill a correspondingly large volume. Various strategies have been developed for reducing the cost, for example by using implicit solvation methods^{38–41}, however it is frequently desirable to treat explicitly the environmental degrees of freedom. Thanks to the fragment approach, the treatment of solvents and other surrounding molecules can readily be achieved in BigDFT with relatively low cost, as we will demonstrate through the example of the fullerene C_{60} in two different environments: when in an aqueous solution and when surrounded by other C_{60} molecules. For each system we constrain a charge of ± 1 to the central C_{60} molecule in order to determine the environmental impact on the ionization potential (IP) and electron affinity (EA).

For traditional DFT calculations with semilocal functionals like LDA it is well known that the above quantities are badly estimated by the frontier orbitals, i.e. the HOMO (LUMO) for the IP (EA). Therefore, in order to extract physically meaningful information, one must either use more expensive beyond-DFT approaches, or instead calculate the IP and EA using the so-called Δ SCF method. This is made possible when explicitly charged calculations are available, i.e. when charged and neutral calculations have energies that can be measured with respect to a common reference. The treatment of the electrostatic potential which is included in the BigDFT code makes such a comparison possible⁴². This latter ap-

proach results in values which match experiment much better than traditional semilocal functionals, indeed our results for the IP and EA of the isolated molecule agree very well with the experimental values of 7.6 eV^{43–45} and 2.7 eV⁴⁶ respectively. In this case, we wish to apply the Δ SCF approach to a molecule in an environment, which, as we will show, is easily achieved using CDFT.


On the other hand, with unconstrained DFT calculations the use of the Δ SCF approach is much more delicate when studying environmental effects with LDA: as the charge tends to be overly delocalized, charged calculations do not simply represent a perturbation from the isolated values, as discussed in more detail below. In other words, the calculated energy differences do not correspond to the IP and EA of C_{60} in an environment, but to a completely different quantity. If one wishes to calculate this quantity it is therefore essential to use CDFT.

1. Computational details

There have been a number of previous studies of C_{60} in water, both experimental and theoretical^{47–53}, however they have mainly focused on neutral fullerenes. Previous research has indicated the existence of a first hydration shell surrounding C_{60} containing between 60 and 65 water molecules^{47–49}, we have therefore chosen to restrict ourselves to systems containing 66 water molecules. We present results for three example structures, which are depicted in Fig. 8(a). They were generated by inserting the C_{60} into water droplets where the water molecules were deposited with random orientations at random positions subject to the room-temperature density of water. The structures were then relaxed until the RMS forces were below 10 meV/Å. For the environment of fullerenes, we limit the cost of the simulations by including six nearest neighbor fullerenes only, so that the system is arranged as a three dimensional cross, as depicted in Fig. 8(b). Each of the fullerenes was considered in its gas-phase structure.

The fragment calculations were performed with a grid spacing of 0.185 Å, while the template calculations were performed using a denser grid of 0.132 Å to ensure accurate reformatting; we used support function radii of 4.23 Å. These values have been chosen such as to ensure the applicability of the Löwdin approach for the weight matrix on the central C_{60} whilst preserving absolute accuracy of the unconstrained calculations to the order of 3 meV/atom, see data in Sec. IIIC2. In order to ensure the support functions are sufficiently accurate for the negatively charged calculations for C_{60} , the template calculation was performed optimizing three additional states (to account for degeneracies).

We continue to use the LDA functional as the difference between the LDA and other treatments like PBE for the IP and EA of fullerenes has previously been shown to be negligible, and reasonable agreement with experiment has also been observed⁵⁴. We also neglected the

FIG. 8. The different environments for C_{60} .

modelling of dispersive terms on the $C_{60} - C_{60}$ interactions due to their negligible impact on frontier orbital eigenvalues and on total energy differences in charged calculations.

2. Testing the fragment approach

The results for the C_{60} structure with a center to center distance of 10 Å are shown in Tab. IV with the corresponding values for the isolated molecule. We also include the cubic scaling results as they allow us to assess the accuracy of the fragment approach for this system. As anticipated, for the isolated molecule the fragment error is of the order of 0.2 eV, which is about 3 meV/atom. In order to confirm that this accuracy is preserved, we also compared unconstrained cubic and fragment calculations for the seven C_{60} structure. To find an unconstrained solution we built the initial guess from the fragment densities in such a manner that the charge was equally distributed among fragments; the final solution remained close to this charge distribution. We found values of -3.681 eV and 6.707 eV, i.e. the difference with the cubic results is of the same magnitude as for the isolated molecule (see Tab. IV).

We have also investigated the effect of varying the separation between the molecules by repeating the constrained fragment and (unconstrained) cubic calculations with center to center separations ranging from 10 Å to 20 Å, which corresponds to a shortest distance between molecules of 3.1 Å to 13.1 Å. The results are plotted in Fig. 9. As expected, for large separations the results tend towards the isolated values. For the unconstrained calculations there is an abrupt change between two different states, whereas with CDFT there is not only a smooth trend, but also an exponential relationship with distance, proving that the fragment approach is sufficiently precise to capture such trends.

Q	isolated		in H_2O			in C_{60} (10 Å)	
	cubic	frag.	A	B	C	cdft	cubic
-1	-2.795	-2.589	-2.017	-2.728	-2.180	-2.854	-3.803
+1	7.648	7.783	7.262	8.033	7.837	7.526	6.685

TABLE IV. Energy differences with neutral, i.e. $E^Q - E^0$ for C_{60} when isolated and in the two environments. Two values are given for the isolated C_{60} : that of the fragment approach, which in this case merely refers to a fixed support function basis as only one fragment is present, and the cubic scaling reference. For the results in water, constrained fragment results are given. For the nearest neighbor results ('in C_{60} '), results are presented for both the constrained fragment and (unconstrained) cubic approaches. The unconstrained results exhibit stronger deviation from the isolated values, showing that the environment is not correctly modeled as it is not acting as a perturbation of the system. Units are in eV.

Thus far, we have only considered calculations with a *shifting* of the template support functions, however we also wish to demonstrate the effectiveness for *rotated* support functions. To this end, for a distance of 10 Å the six outer fullerenes were collectively rotated along the z -axis by angles of 15°, 45° and 90°, with the orientation of the central molecule remaining unchanged. This was found to have a negligible impact on the IP and EA, with a difference in the values for the various orientations of around 0.01 eV for the constrained fragment calculations compared with 0.05 eV for the unconstrained cubic calculations. Such values are too small to be significant compared to the errors associated with the basis. In order to determine whether the energies are truly unaffected by the orientation, it would be necessary to account for the dispersion effects which are not captured by the LDA. However, the fact that no spurious errors are introduced by the rotating of the support functions serves to further confirm the accuracy of the reformatting scheme.




FIG. 9. Variation in electron affinity and ionization potential for increasing separations, where the distance is measured between the centers of neighboring molecules. The energy plotted is relative to the isolated value in the respective basis, i.e. $\Delta E = (E_{\text{isol.}}^Q - E_{\text{isol.}}^0) - (E_{\text{full}}^Q - E_{\text{full}}^0)$, where Q indicates the charge state, ‘full’ refers to the 7 molecule system and ‘isol.’ refers to the isolated molecule.

3. Comparison of environments

We now return to the results in Tab. IV in order to compare the effect of the two environments. As expected, for both environments the constrained values remain relatively close to the isolated results, certainly much closer than the unconstrained results for the fullerene environment. In a sense, when the constraint is enforced, the presence of the surrounding molecules could be thought of as a perturbation on the isolated state, although the strength of the perturbation is clearly much stronger for the water. To further explore this, we have also plotted differences in the converged electronic densities between the neutral and charged calculations in Fig. 10. The effects of the charge constraint are clearly visible, with the excess charge distributed across the molecules for the cubic calculation and much more clearly localized for the constrained calculations. Furthermore, the charge difference on the central molecule for the constrained calculations clearly retains the same character as the respective isolated density, with the excess or deficit of charge also resulting in an induced dipole on the neighboring molecules. As expected, the impact of the water is stronger than the neighboring fullerenes, where the closer proximity and stronger dipole moment of the water molecules results in a stronger deviation from the isolated density difference. Similar behavior has also been observed for the water structures which are not depicted.

As can be seen from the variation in IP and EA for structures A, B and C (Tab. IV), the effect of the water is not only stronger than that of the neighboring fullerenes, but the dependence on the structure is also quite significant. Furthermore, we have also observed that the resulting energies are strongly affected by the choice of

weight function, which is not the case for the fullerene environment. There is some freedom in the procedure for optimizing the template support functions, e.g. the localization radii and the number of additional states included; we tested a few of the different options. For the fullerene environment the variation in calculated IP and EA due to the choice of template parameters were small and systematic, whereas for the aqueous environment the variation was much stronger. Indeed, the fragment approach provides an ideal setup to explore the impact of different weight functions – such a strong dependence would be harder to detect when considering only two or three different choices. In future the choice of weight function could also be decoupled from the fragment basis to allow for a more thorough exploration of its influence in constraining the charge.

Of course, in order to correctly assess the impact of the environment, one should go beyond the model structures used here, both in terms of the size and procedure used to generate them. Furthermore, in the case of the water, proper sampling should be performed over a number of different configurations, which should be generated at the correct temperature e.g. using molecular dynamics^{50,52} or Monte Carlo^{47,48} simulations. However, aside from the generation of input structures and any eventual relaxations of the atomic coordinates, the fragment calculations are quick and easy to perform, requiring little additional setup aside from the template calculations. Furthermore, we have demonstrated both the accuracy and flexibility of the fragment approach for such systems. As such, given appropriate atomic coordinates this work could easily be extended in future to a large number of configurations for both environments, or indeed applied to other fullerenes or solvents.

IV. CONCLUSION

We have presented a method for constrained DFT calculations on large systems, using a fragment based scheme. This has been implemented in the BigDFT electronic structure code within a recently developed framework, which uses a basis of localized support functions represented in an underlying wavelet grid to achieve linear scaling behavior with respect to system size while retaining the systematic accuracy of the underlying grid. The division of a given system into fragments (ideally distinct molecules), each with its own associated support functions leads to a natural approach to CDFT, where the charge is constrained to a given fragment via a Löwdin like definition of the weight function. This Löwdin approach can also be used to straightforwardly calculate atomic charges, as we have demonstrated.

Furthermore, by using a reformatting scheme which enables the reuse of support functions for identical fragments, irrespective of their position or orientation in the system, we are able to further reduce the cost of simulations by an order of magnitude, as the support functions




FIG. 10. Density differences between the neutral and charged calculations for the fullerene when isolated, when surrounded by other fullerenes (with a center to center distance of 10 Å) and when in water (structure A). The densities are plotted on the central plane with a logscale, with red (blue) indicating an increase (decrease) in the electronic charge density with respect to the neutral.

can be separately optimized for each ‘template’ fragment and used as a fixed basis in the system of interest. The properties of the wavelet basis set ease the implementation of reformatting a numerical field given in real space, so that we were able to implement this scheme in a manner which is both efficient and accurate. The flexibility of this method, together with the ability of the BigDFT code to treat systems with many atoms (see e.g. Ref.²⁵), makes it ideally suited for both neutral and charged calculations on very large systems at modest computational cost.

We have presented results from two previously studied systems in order to validate our approach, as well as an example application. For this latter point, we have performed calculations on C_{60} in two different environments, namely a model nearest neighbor system containing seven fullerenes and in an aqueous solution. The effects of the constraint are clearly visible in the electronic densities, which we have compared to the unconstrained and isolated results. We have also shown that the presence of water has both a stronger impact on the results and a stronger dependence on the choice of weight function than the presence of neighboring fullerenes.

The reformatting approach described here has another key benefit aside from reducing the cost of such calcu-

lations: as the basis set for each fragment will remain equivalent following the reformatting, the computational setup provided by our approach is ideal where Hamiltonian matrix elements of the whole system have to be considered. For example, for electronic coupling matrix elements (‘transfer integrals’) between two identical monomers, the basis set for each monomer will remain equivalent following the reformatting, so that there is no ambiguity in the sign of the coupling matrix elements. In contrast, for support functions optimized from scratch there is no guarantee that the phase for both the support functions and wavefunction coefficients will be identical between the two molecules and thus the sign of the transfer integral cannot be determined. Indeed, we will be publishing results in the near future for such an application based on the framework presented in the current work⁵⁵.

In the future we also hope to extend this work to permit calculations on realistic nanoscale devices, using a multi-scale approach. In the first instance this would involve changing the definition of a fragment to an individual atom, which naturally leads to a DFT based tight-binding like method. This formalism also allows the correct definition of an ‘embedded’ approach, where different regions of a simulation cell are treated at differ-

ent levels of precision, e.g. for a point defect in a bulk semiconductor, with higher accuracy close to the defect. Work is ongoing in this direction.

ACKNOWLEDGEMENTS

We acknowledge funding from the European project MMM@HPC (RI-261594), the CEA-NANOSCIENCE

BigPOL project and the ANR projects SAMSON (ANR-AA08-COSI-015) and NEWCASTLE (ANR-2010-COSI-005-01). This research used resources of the Argonne Leadership Computing Facility at Argonne National Laboratory, which is supported by the Office of Science of the U.S. Department of Energy under contract DE-AC02-06CH11357. CPU time was also provided by IDRIS (project i2014096905).

* lratcliff@anl.gov

- ¹ P. Hohenberg and W. Kohn, Phys. Rev. **136**, B864 (1964)
- ² W. Kohn and L. J. Sham, Phys. Rev. **140**, A1133 (1965)
- ³ J. P. Perdew and A. Zunger, Phys. Rev. B **23**, 5048 (1981)
- ⁴ Y. Zhang and W. Yang, J. Chem. Phys. **109**, 2604 (1998)
- ⁵ P. H. Dederichs, S. Blügel, R. Zeller, and H. Akai, Phys. Rev. Lett. **53**, 2512 (1984)
- ⁶ E. Runge and E. K. U. Gross, Phys. Rev. Lett. **52**, 997 (1984)
- ⁷ A. M. P. Sena, T. Miyazaki, and D. R. Bowler, J. Chem. Theory Comput. **7**, 884 (2011)
- ⁸ H. Oberhofer and J. Blumberger, J. Chem. Phys. **131**, 064101 (2009)
- ⁹ J. Řezáč, B. Lévy, I. Demachy, and A. de la Lande, J. Chem. Theory Comput. **8**, 418 (2012)
- ¹⁰ A. M. Souza, I. Rungger, C. D. Pemmaraju, U. Schwingenschloegl, and S. Sanvito, Phys. Rev. B **88**, 165112 (2013)
- ¹¹ Q. Wu and T. Van Voorhis, J. Chem. Phys. **125**, 164105 (2006)
- ¹² H. Oberhofer and J. Blumberger, J. Chem. Phys. **133**, 244105 (2010)
- ¹³ B. Kaduk, T. Kowalczyk, and T. Van Voorhis, Chem. Rev. **112**, 321 (2012)
- ¹⁴ I. Daubechies, *Ten Lectures on Wavelets* (SIAM, 1992)
- ¹⁵ S. Mohr, L. E. Ratcliff, P. Boulanger, L. Genovese, D. Caliste, T. Deutsch, and S. Goedecker, J. Chem. Phys. **140**, 204110 (2014)
- ¹⁶ L. Genovese, A. Neelov, S. Goedecker, T. Deutsch, S. A. Ghasemi, A. Willand, D. Caliste, O. Zilberberg, M. Rayson, A. Bergman, and R. Schneider, J. Chem. Phys. **129**, 014109 (2008)
- ¹⁷ W. Kohn, Phys. Rev. Lett. **76**, 3168 (1996)
- ¹⁸ E. Prodan and W. Kohn, PNAS **102**, 11635 (2005)
- ¹⁹ S. Ismail-Beigi and T. A. Arias, Phys. Rev. Lett. **82**, 2127 (1999)
- ²⁰ W. Kohn, Int. J. Quantum Chem. **56**, 229 (1995)
- ²¹ L. He and D. Vanderbilt, Phys. Rev. Lett. **86**, 5341 (2001)
- ²² R. Baer and M. Head-Gordon, Phys. Rev. Lett. **79**, 3962 (1997)
- ²³ C.-K. Skylaris, P. D. Haynes, A. A. Mostofi, and M. C. Payne, J. Chem. Phys. **122**, 084119 (2005)
- ²⁴ D. R. Bowler and T. Miyazaki, J. Phys.: Condens. Matter **22**, 074207 (2010)
- ²⁵ S. Mohr, L. E. Ratcliff, L. Genovese, D. Caliste, P. Boulanger, S. Goedecker, and T. Deutsch, arXiv:1501.05884(submitted)
- ²⁶ Q. Wu and T. Van Voorhis, Phys. Rev. A **72**, 024502 (2005)
- ²⁷ K. Senthilkumar, F. C. Grozema, F. M. Bickelhaupt, and L. D. A. Siebbeles, J. Chem. Phys. **119**, 9809 (2003)
- ²⁸ K. Senthilkumar, F. C. Grozema, C. F. Guerra, F. M. Bickelhaupt, F. D. Lewis, Y. A. Berlin, M. A. Ratner, and L. D. A. Siebbeles, J. Am. Chem. Soc. **127**, 14894 (2005)
- ²⁹ W. Kabsch, Acta Crystallogr. A **34**, 827 (Sep 1978)
- ³⁰ G. Wahba, SIAM Rev. **7**, 409 (1965)
- ³¹ F. L. Markley, J. Astronaut. Sci. **36**, 245 (1988)
- ³² A. Neelov and S. Goedecker, J. Comp. Phys. **217**, 312 (2006)
- ³³ D. M. Ceperley and B. J. Alder, Phys. Rev. Lett. **45**, 566 (1980)
- ³⁴ C. Hartwigsen, S. Goedecker, and J. Hutter, Phys. Rev. B **58**, 3641 (1998)
- ³⁵ Q. Wu and T. Van Voorhis, J. Chem. Theory Comput. **2**, 765 (2006)
- ³⁶ I. Duchemin, T. Deutsch, and X. Blase, Phys. Rev. Lett. **109**, 167801 (2012)
- ³⁷ A. Dreuw and M. Head-Gordon, J. Am. Chem. Soc. **126**, 4007 (2004)
- ³⁸ J. Tomasi and M. Persico, Chem. Rev. **94**, 2027 (1994)
- ³⁹ D. A. Scherlis, J.-L. Fattebert, F. Gygi, M. Cococcioni, and N. Marzari, J. Chem. Phys. **124**, 074103 (2006)
- ⁴⁰ J. Fosso-Tande and R. J. Harrison, Chem. Phys. Lett. **561-562**, 179 (2013)
- ⁴¹ J. Dziedzic, H. H. Helal, C.-K. Skylaris, A. A. Mostofi, and M. C. Payne, Europhys. Lett. **95**, 43001 (2011)
- ⁴² A. Cerioni, L. Genovese, A. Mirone, and V. A. Sole, J. Chem. Phys. **137**, 134108 (2012)
- ⁴³ D. Lichtenberger, K. Nebesny, C. Ray, D. Huffman, and L. Lamb, Chem. Phys. Lett. **176**, 203 (1991)
- ⁴⁴ J. Devries, H. Steger, B. Kamke, C. Menzel, B. Weisser, W. Kamke, and I. Hertel, Chem. Phys. Lett. **188**, 159 (1992)
- ⁴⁵ J. A. Zimmerman, J. R. Eyler, S. B. H. Bach, and S. W. McElvany, J. Chem. Phys. **94**, 3556 (1991)
- ⁴⁶ X.-B. Wang, C.-F. Ding, and L.-S. Wang, J. Chem. Phys. **110**, 8217 (1999)
- ⁴⁷ R. Rivelino and F. de Brito Mota, Nano Lett. **7**, 1526 (2007)
- ⁴⁸ R. Rivelino, A. M. Maniero, F. V. Prudente, and L. S. Costa, Carbon **44**, 2925 (2006)
- ⁴⁹ P. Scharff, K. Risch, L. Carta-Abelmann, I. Dmytruk, M. Bilyi, O. Golub, A. Khavryuchenko, E. Buzaneva, V. Aksenov, M. Avdeev, Y. Prylutsky, and S. Durov, Carbon **42**, 1203 (2004)
- ⁵⁰ L. Li, D. Bedrov, and G. D. Smith, J. Chem. Phys. **123**, 204504 (2005)
- ⁵¹ Y. I. Prylutsky, A. S. Buchelnikov, D. P. Voronin, V. V. Kostjukov, U. Ritter, J. A. Parkinson, and M. P. Evstigneev, Phys. Chem. Chem. Phys. **15**, 9351 (2013)
- ⁵² S. Banerjee, J. Chem. Phys. **138**, 044318 (2013)

- ⁵³ N. Choudhury, J. Chem. Phys. **125**, 034502 (2006)
- ⁵⁴ M. L. Tiago, P. R. C. Kent, R. Q. Hood, and F. A. Reboredo, J. Chem. Phys. **129**, 084311 (2008)
- ⁵⁵ L. E. Ratcliff, L. Grisanti, L. Genovese, T. Deutsch, T. Neumann, D. Danilov, W. Wenzel, D. Beljonne, and

J. Cornil(submitted)